Distributed averaging in the presence of a sparse cut

Hariharan Narayanan
Department of Computer Science,
University of Chicago

May 2, 2008

Abstract

We consider the question of averaging on a graph that has one sparse cut separating two subgraphs that are internally well connected. While there has been a large body of work devoted to algorithms for distributed averaging, nearly all algorithms involve only *convex* updates. In this paper, we suggest that *non-convex* updates can lead to significant improvements. We do so by exhibiting a decentralized algorithm for graphs with one sparse cut that uses non-convex averages and has an averaging time that can be significantly smaller than the averaging time of known distributed algorithms, such as those of [3, 2]. We use stochastic dominance to prove this result in a way that may be of independent interest.

1 Introduction

Consider a Graph $G = (V, E)$, where i.i.d Poisson clocks with rate 1 are associated with each edge[1]. We represent the "true" real valued time by T. Each node v_i holds a value $x_i(T)$ at time T. Let the average value held by the nodes be x_{av}. Every time an edge $e = (v, w)$ ticks, it updates the values of vertices adjacent to it on the basis of present and past values of v, w and their immediate neighbors according to some algorithm $\mathcal{A}$. There is an extensive body of work surrounding the subject of gossip algorithms in various contexts. Non-convex updates have been used in the context of a second order diffusion for load balancing [5] in a slightly different setting. The idea there was to take into account the value of the nodes during the previous two time steps rather than just the previous one, (in a synchronous setting), and set the future value of a node to a non-convex linear combination of the past values of some

[1]This model can be simulated using previous models such as [2] by allocating edges to nodes and equipping nodes with multiple i.i.d poisson clocks.

of its neighbors. There is also a line of research on averaging algorithms having two time scales, [1, 4] which is closely related to the present paper.

In a previous paper [6], we considered the use of non-convex combinations for gossip on a geographic random graph on n nodes. There we showed that one can achieve averaging using $n^{1+o(1)}$ updates if one is willing to allow a certain amount of centralized control.The main technical difficulty in using non-convex updates is that they can skew the values held by nodes in the short term. We show that nonetheless, in the long term this leads to faster averaging. Let the values held by the nodes by $X(T) = (x_1(T), \ldots, x_{|V|}(T))^T$. We study distributed averaging algorithms $\mathcal{A}$ which result in

$$\lim_{T \to \infty} X(T) = x_{av}\mathbf{1},$$

where x_{av} is invariant under the passage of time. and show that in some cases there is an exponential speed-up in n if one allows the use of non-convex updates, as opposed to only convex ones.

Definition 1 *Let*

$$\operatorname{var} X(t) := \frac{\sum_{i=1}^{|V|}(x_i(t) - x_{av})^2}{|V|}.$$

Let

$$T_{av} = \sup_{x \in \mathbb{R}^{|V|}} \inf_t \ \mathbb{P}\left[\exists T > t,\ \frac{\operatorname{var} X(T)}{\operatorname{var} X(0)} > \frac{1}{e^2} \ \middle|\ X(0) = x \right] < \frac{1}{e}.$$

Notation 1 *Let a connected graph $G = (V, E)$ have a partition into connected graphs $G_1 = (V_1, E_1)$, and $G_2 = (V_2, E_2)$. Specifically, every vertex in V is either in V_1 or V_2, and every edge in E belongs to either E_1 or to E_2, or to the set of edges E_{12} that have one endpoint in V_1 and one in V_2. Let $|V_1| = n_1$, $|V_2| = n_2$ where without loss of generality, $n_1 \leq n_2$ and $|V| = n$. Let $T_{van}(G_1)$ and $T_{van}(G_2)$ be the averaging times of the "vanilla" algorithm that replaces at the clock tick of an edge e the values of the endpoints of e by the arithmetic mean of the two, applied to G_1 and G_2 respectively.*

Definition 2 *Let* $\mathbf{C}$ *denote the set of algorithms that use only convex updates of the form*

1. $x_i(t^+) = \alpha x_i(t^-) + \beta x_j(t^-)$.
2. $x_j(t^+) = \alpha x_j(t^-) + \beta x_i(t^-)$.

where $\alpha \in [0, 1]$ and $\alpha + \beta = 1$.

These updates have been extensively studied, see for example [3, 2].

Theorem 1 *The averaging time of any distributed algorithm in* $\mathbf{C}$ *is* $\Omega(\frac{\min(|V_1|,|V_2|)}{|E_{12}|})$

Theorem 2 *The averaging time of* $\mathcal{A}$ *is* $O(\log n(T_{van}(G_1) + T_{van}(G_2)))$.

Note that in the case where G_1 and G_2 are sufficiently well connected internally but poorly connected to each other, $\mathcal{A}$ outperforms any algorithm in $\mathbf{C}$. In fact for the graph G' obtained by joining two complete graphs G_1', G_2' each having $\frac{n}{2}$ vertices by a single edge, $\Omega(\frac{\min(|V_1'|,|V_2'|)}{|E_{12}'|}) = \Omega(n)$, while $O(\log n(T_{av}(G_1') + T_{av}(G_2'))) = O(\log n)$.

1.0.1 Algorithm $\mathcal{A}$

Let the vertices of G_1 be labeled by $[n_1]$ and those of G_2 by $[n_2]\setminus[n_1]$, where $[n] := \{1, \dots, n\}$. Let $e_c = (v_{n_1}, v_{n_1+1}$ be a fixed edge belonging to E_{12}. Let the time of the k^{th} clock tick of an edge e be t. Let $C >> 1$ be a sufficiently large absolute constant (independent of n.)

- If the edge e is $e_c = (v_{n_1}, v_{n_1+1})$,
  1. If $k \equiv -1 \mod (\lceil C(T_{van}(G_1) + T_{van}(G_2)) \ln n \rceil)$
     (a) $x_{n_1}(t^+) = x_{n_1}(t^-) + n_1 \{x_{n_1+1}(t^-) - x_{n_1}(t^-)\}$
     (b) $x_{n_1+1}(t^+) = x_{n_1+1}(t^-) - n_1 \{x_{n_1+1}(t^-) - x_{n_1}(t^-)\}$
  2. If $k \not\equiv -1 \mod (\lceil C(T_{van}(G_1) + T_{van}(G_2)) \ln n \rceil)$ make no update.
- If the edge e is $(v_i, v_j) \notin E_{12}$
  1. $x_i(t^+) = \frac{x_i(t^-)+x_j(t^-)}{2}$.
  2. $x_j(t^+) = \frac{x_i(t^-)+x_j(t^-)}{2}$.
- If $e \in E_{12} \setminus \{e_c\}$ make no update.

---

2 Limitations of convex combinations

Given a function $a(t)$, let its right limit at t be denoted by $a(t+)$ and its left limit at t by $a(t^-)$. Consider an algorithm $\mathcal{C} \in \mathbf{C}$.
Let us consider the initial condition where $X(0)$ is the vector that is 1 on vertices $v_1, \dots, v_{n_1}$ of G_1 and $-\frac{n_1}{n_2}$ on vertices $v_{n_1+1}, \dots, v_n$ of G_2. Let us denote $\frac{\sum_{i=1}^{n_1} x_i(t)}{n_1}$ by $y(t)$ and $\frac{\sum_{i=n_1+1}^{n_2} x_i(t)}{n_2}$ by $z(t)$. In the model we have considered, with probability 1, at no time does more than one

clock tick.
In the course of the execution any algorithm in **C** $y(t)$ can change only during clock ticks of e_c and the same holds for $z(t)$. This is because during a clock tick of any other edge, both of whose end-vertices lie in G_1 or in G_2, $y(t)$ and $z(t)$ do not change. The vertices adjacent to e_c can change by at most 2 across these instants. Further, the values $x_n(t)$ and $x_{n+1}(t)$ are seen to lie in the interval

$$[\min_{i\in|V|} x_i(0), \max_{i\in|V|} x_i(0)] \subseteq [-1, 1].$$

If the clock of e_c ticks at time t, we therefore find that

$$|y(t^+) - y(t^-)| \leq \frac{2}{n_1}, \tag{1}$$

The number of clocks ticks of e_c until time t is a Poisson random variable whose mean is t.

A direct calculation tells us that

$$\mathrm{var}(X(t)) \quad \geq \quad \frac{n_1 y(t)^2}{n}. \tag{2}$$

To obtain a lower bound for $y(t)^2$, we note that the total number of times the clocks of edges belonging to E_{12} tick is a Poisson random variable ν_t with mean $t|E_{12}|$. It follows from Inequality (1) that $y(t) \geq 1 - \frac{2\nu_t}{n_1}$.

$$\begin{aligned}
|E_{12}|T_{av} &= \mathbb{E}[\nu_{T_{av}}] \\
&\geq \mathbb{P}\left[\nu_{T_{av}} \geq (1-\frac{1}{e})\frac{n_1}{4}\right](1-\frac{1}{e})\frac{n_1}{4}
\end{aligned}$$

However

$$\mathbb{P}\left[\nu_{T_{av}} \geq (1-\frac{1}{e})\frac{n_1}{4}\right]$$

must be large, because otherwise $y(T_{av})$ would probably be large. More precisely,

$$\begin{aligned}
\mathbb{P}\left[\nu_{T_{av}} \geq (1-\frac{1}{e})\frac{n_1}{4}\right] &\geq 1 - \mathbb{P}\left[\ \exists T > T_{av},\ \mathrm{var}\, X(T) > \frac{1}{e^2}\ \right] \\
&\geq 1 - \frac{1}{e}
\end{aligned}$$

Therefore,

$$\begin{aligned}
T_{av} &\geq \mathbb{P}\left[\nu_{T_{av}} \geq (1-\frac{1}{e})\frac{n_1}{4|E_{12}|}\right](1-\frac{1}{e})\frac{n_1}{4} \\
&\geq \Omega(\frac{n_1}{|E_{12}|})
\end{aligned}$$

3 Using non-convex combinations

3.0.2 Analysis

Since T_{av} is defined in terms of variance and algorithm $\mathcal{A}$ uses only linear updates, we may subtract out the mean from each $X_i(0)$ and it is sufficient to analyze the case when $x_{av} = 0$. Let $V_1 = [n_1]$ and $V_2 = [n] \setminus [n_1]$. Let $\mu_1(t) = \frac{\sum_{i=1}^{n_1} x_i(t)}{n_1}$ and $\mu_2 = \frac{\sum_{i=n_1}^{2n} x_i(t)}{n}$ and $\mu(t) = |\mu_1(t)| + |\mu_2(t)|$. Let

$$\sigma(t) = \sqrt{\frac{\sum_{i=1}^{n_1}(x_i(t) - \mu_1(t))^2 + \sum_{n_1+1}^{n}(x_i(t) - \mu_2(t))^2}{n}}.$$

We consider time instants $T_1, T_2, \dots$ where T_i is the instant at which the clock of edge e ticks for the $\lceil iC(T_{van}(G_1) + T_{van}(G_2)) \ln n \rceil^{th}$ time. Observe that the value of $\mu(t)$ changes only across time instants $T_k, k = 1, 2, \dots$.
The amount by which $x_{n_1}(t)$ and $x_{n_1+1}(t)$ deviate from $\mu_1(t)$ and $\mu_2(t)$ respectively, can be seen to be bounded above by $\sqrt{n}\sigma(t)$

$$\max\left\{|x_{n_1}(t) - \mu_1(t)|\ , |x_{n+1}(t) - \mu_2(t)|\right\} \ \leq \ \sqrt{n}\sigma(t). \tag{3}$$

We now examine the evolution of $\sigma(T_k^+)$ and $\mu(T_k^+)$ as $k \to \infty$. The statements below are true if C is a sufficiently large universal constant (independent of n).

From T_k^+ to T_{k+1}^-, independent of x,

$$\mathbb{P}\left[\sigma(T_{k+1}^-) \geq \frac{\sigma(T_k^+)}{n^6} \,\middle|\, X(T_k^+) = x\right] \ \leq \ \frac{1}{4n} \tag{4}$$

$$\mu(T_{k+1}^-) \ = \ \mu(T_k^+). \tag{5}$$

Because of inequality (3), from T_k^+ to T_{k+1}^-

$$\sigma(T_{k+1}^+) \leq n(\sigma(T_{k+1}^-) + |\mu(T_{k+1}^-)|) \tag{6}$$

$$|\mu(T_{k+1}^+)| \leq n^{\frac{3}{2}}\ \sigma(T_{k+1}^-) \tag{7}$$

$$\text{var } X(t) = \mu(t)^2 + \sigma(t)^2.$$

We deduce from the above that

$$\mathbb{P}\left[\text{var } X(T_{k+1}^+) \geq \frac{\text{var } X(T_k^+)}{n^4}\right] \ \leq \ \frac{1}{4n} \tag{8}$$

Let A_k be the (random) operator obtained by composing the linear updates from time T_k^+ to T_{k+1}^+. Let $\|A\|$ denote the norm of an operator acting from ℓ_2 to ℓ_2

$$\|A\| = \frac{\sup_{x\in\mathbb{R}^n} \|Ax\|_2}{\|x\|_2}.$$

Lemma 1

$$\mathbb{P}\left[\|A_k\|^2 \geq \frac{1}{n^3}\right] \;\leq\; \frac{1}{2} \tag{9}$$

To see this, let $v_1, \ldots, v_n$ be the canonical basis for $\mathbb{R}^n$. For any unit vector

$$x = \sum_{i=1}^n \lambda_i v_i$$

Then,

$$\|A_k(x)\| \;\leq\; \sum_{i=1}^n |\lambda_i|\; \|A_k(v_i)\| \quad \text{(Triangle Inequality)} \tag{10}$$

$$\leq\; \sqrt{\sum_{i=1}^n \|A_k(v_i)\|^2} \quad \text{(Cauchy-Schwartz inequality)} \tag{11}$$

The Lemma now follows from Inequality (8) by an application of the Union Bound. □
Moreover, we observe by construction that the norm of A_k is less or equal to n,

$$\|A_k\| \;\leq\; n \tag{12}$$

Note that $\log(\mathrm{var}\, X(T_k^+))$ defines a random process (that is *not* Markov). The updates A_k from time T_k^+ to T_{k+1}^+ for successive k are i.i.d random operators acting on $\mathbb{R}^{2n}$. Note that

$$\log(\mathrm{var}\, X(T_k^+)) - \log(\mathrm{var}\, X(0)) \leq \sum_{i=1}^k \log \|A_i\|$$

due to the presence of the supremum in the definition of operator norm.

$$W_k := \sum_{i=1}^k \log \|A_i\|$$

is a random walk on the real line for $k = 1, \ldots, \infty$.

The last and perhaps most important ingredient is that of *stochastic dominance*. It follows from Lemma 1 and Equation 12 that the random walk $\{W_k\}$ can be coupled with a random

walk $\{\tilde{W}_k\}$ that is always to the right of it on the real line, i. e. for all k, $W_k \leq \tilde{W}_k$, where the increments

$$\begin{aligned} \tilde{W}_{k+1} - \tilde{W}_k &= \log n \quad (\text{with probability } \frac{1}{2}) && (13) \\ &= -\frac{3}{2}\log n \quad (\text{with probability } \frac{1}{2}.) && (14) \end{aligned}$$

Noting that by construction,

$$\log(\operatorname{var} X(T_k^+)) - \log(\operatorname{var} X(0)) \leq \tilde{W}_k, \tag{15}$$

it follows that T_{av} is upper bounded by any t_0 which satisfies

$$\mathbb{P}\Big[\ \forall T > t_0,\ \tilde{W}_T \leq -2\Big] > 1 - \frac{1}{e}.$$

Note that $\mathbb{E}[\tilde{W}_k] = -\frac{k \log n}{2}$ and $\mathbb{E}[\operatorname{var}\tilde{W}_k] = \frac{9k}{16}\log^2 n$.

In order to proceed, we shall need the following inequality about simple unbiased random walk $\{S_k\}_{k\geq 0}$ on $\mathbb{Z}$ starting at 0.

Theorem 3 *There exist constants c, β such that for any $n \in \mathbb{Z}$, $s > 0$*

$$\mathbb{P}[S_n \geq s\sqrt{n}] \leq ce^{-\beta s^2}.$$

Using this fact,

$$\mathbb{P}[\forall T > t_0,\ \tilde{W}_T \leq -2] = \mathbb{P}[\forall T > t_0, (\log n)(S_T - \frac{T}{2}) \leq -2] \tag{16}$$

For large n, this is the same as

$$\mathbb{P}[\forall T > t_0, S_T < \frac{T}{2}] \geq 1 - \sum_{T > t_0} ce^{-\beta T/4}.$$

Clearly, there is a constant t_0 independent of n such that $1 - \sum_{T>t_0} ce^{-\beta T/4} > 1 - \frac{1}{e}$. This completes the proof. □

4 Acknowledgement

I am grateful to Dimitris Achlioptas, Vivek Borkar, Stephen Boyd and Steven Lalley for many helpful discussions.

References

[1] V. S. Borkar.
Stochastic approximation with two time-scales. *Systems and Control letters.* (1997)

[2] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah.
Gossip algorithms : Design, analysis and applications.
In *Proceedings of the 24th Conference of the IEEE Communications Society (INFOCOM 2005)*, 2005.

[3] D. Bertsekas, J. Tsitsiklis.
Parallel and Distributed Computation: Numerical Methods, Prentice-Hall, 1989

[4] V. Konda and J. Tsitsiklis.
Convergence rate of a linear two time scale stochastic approximation. *Annals of Applied Probability* 2004.

[5] S. Muthukrishnan, B. Ghosh and M. H. Schultz
First and Second-Order Diffusive Methods for Rapid, Coarse, Distributed Load Balancing.
Theory of Computing Systems, Volume 31, Number 4, December, 1998

[6] H. Narayanan.
Geographic Gossip on Geometric Random Graphs via Affine Combinations.
In *Principles of Distributed Computing (PODC)*, 2007